

FORM PTO-1390 (REV 11-2000)	U.S. DEPARTMENT OF COMMERCE PATENT AND TRADEMARK OFFICE	ATTORNEY'S DOCKET NUMBER 36-1536
TRANSMITTAL LETTER TO THE UNITED STATES DESIGNATED/ELECTED OFFICE (DO/EO/US) CONCERNING A FILING UNDER 35 U.S.C. 371		U.S. APPLICATION NO. (If known, see 37 C.F.R. 1.5) 10/069359 Unknown
INTERNATIONAL APPLICATION NO. PCT/GB00/03684	INTERNATIONAL FILING DATE 25 September 2000	PRIORITY DATE CLAIMED 24 September 1999

TITLE OF INVENTION

PACKET NETWORK INTERFACING

APPLICANT(S) FOR DO/EO/US


HOVELL et al

Applicant herewith submits to the United States Designated/Elected Office (DO/EO/US) the following items and other information:

1. ☒ This is a **FIRST** submission of items concerning a filing under 35 U.S.C. 371.
2. ☐ This is a **SECOND** or **SUBSEQUENT** submission of items concerning a filing under 35 U.S.C. 371.
3. ☒ This is an express request to begin national examination procedures (35 U.S.C. 371(f)). The submission must include items (5), (6), (9) and (21) indicated below.
4. ☒ The U.S. has been elected by the expiration of 19 months from the priority date (Article 31).
A copy of the International Application as filed (35 U.S.C. 371(c)(2)).
 - a. ☒ is attached hereto (required only if not communicated by the International Bureau).
 - b. ☒ has been communicated by the International Bureau.
 - c. ☐ is not required, as the application was filed in the United States Receiving Office (RO/US).
5. ☐ An English language translation of the International Application as filed (35 U.S.C. 371(c)(2)).
 - a. ☐ is attached hereto.
 - b. ☐ has been previously submitted under 35 U.S.C. 154(d)(4).
6. ☐ Amendments to the claims of the International Application under PCT Article 19 (35 U.S.C. 371(c)(3))
 - a. ☐ are attached hereto (required only if not communicated by the International Bureau).
 - b. ☐ have been communicated by the International Bureau.
 - c. ☐ have not been made; however, the time limit for making such amendments has **NOT** expired.
 - d. ☐ have not been made and will not be made.
8. ☐ An English language translation of the amendments to the claims under PCT Article 19 (35 U.S.C. 371(c)(3)).
9. ☒ An oath or declaration of the inventor(s) (35 U.S.C. 371(c)(4)).
10. ☐ A English language translation of the annexes of the International Preliminary Examination Report under PCT Article 36 (35 U.S.C. 371(c)(5)).

Items 11 To 20 below concern document(s) or information included:

11. ☐ An Information Disclosure Statement under 37 C.F.R. 1.97 and 1.98.
12. ☒ An assignment document for recording. A separate cover sheet in compliance with 37 C.F.R. 3.28 and 3.31 is included.
13. ☒ A **FIRST** preliminary amendment.
14. ☐ A **SECOND** or **SUBSEQUENT** preliminary amendment.
15. ☐ A substitute specification.
16. ☐ A change of power of attorney and/or address letter.
17. ☐ A computer-readable form of the sequence listing in accordance with PCT Rule 13ter.2 and 35 U.S.C. 1.821-1.825.
18. ☐ A second copy of the published international application under 35 U.S.C. 154(d)(4).
19. ☐ A second copy of the English language translation of the international application under 35 U.S.C. 154(d)(4).
20. ☒ Other items or information. AMENDED SHEETS -- pages 2-5, 5a, 5b, 6, 6a and 18 through 23 (claims 1-21)

U.S. APPLICATION NO. (If known, see 37 C.F.R. 1.51) Unknown 069359		INTERNATIONAL APPLICATION NO. PCT/GB00/03684		ATTORNEY'S DOCKET NUMBER 36-1536																															
21. <input checked="" type="checkbox"/> The following fees are submitted: BASIC NATIONAL FEE (37 C.F.R. 1.492(a)(1)-(5)): -- Neither international preliminary examination fee (37 C.F.R. 1.482) nor international search fee (37 C.F.R. 1.445(a)(2)) paid to USPTO and International Search Report not prepared by the EPO or JPO\$1040.00 -- International preliminary examination fee (37 C.F.R. 1.482) not paid to USPTO but International Search Report prepared by the EPO or JPO\$890.00 -- International preliminary examination fee (37 C.F.R. 1.482) not paid to USPTO but international search fee (37 C.F.R. 1.445(a)(2)) paid to USPTO\$740.00 -- International preliminary examination fee (37 C.F.R. 1.482) paid to USPTO but all claims did not satisfy provisions of PCT Article 33(1)-(4)\$710.00 -- International preliminary examination fee (37 C.F.R. 1.482) paid to USPTO and all claims satisfied provisions of PCT Article 33(1)-(4)\$100.00 ENTER APPROPRIATE BASIC FEE AMOUNT =				CALCULATIONS PTO USE ONLY \$ 890.00 \$ 0.00																															
Surcharge of \$130.00 for furnishing the oath or declaration later than <input type="checkbox"/> 20 <input type="checkbox"/> 30 months from the earliest claimed priority date (37 C.F.R. 1.492(e)). <table border="1" style="width:100%; border-collapse: collapse;"> <thead> <tr> <th>CLAIMS</th> <th>NUMBER FILED</th> <th>NUMBER EXTRA</th> <th>RATE</th> <th></th> <th></th> </tr> </thead> <tbody> <tr> <td>Total Claims</td> <td>19</td> <td>-20 =</td> <td>0</td> <td>X</td> <td>\$18.00</td> </tr> <tr> <td>Independent Claims</td> <td>4</td> <td>-3 =</td> <td>1</td> <td>X</td> <td>\$84.00</td> </tr> <tr> <td colspan="4">MULTIPLE DEPENDENT CLAIMS(S) (if applicable)</td> <td></td> <td>\$280.00</td> </tr> <tr> <td colspan="4">TOTAL OF ABOVE CALCULATIONS =</td> <td></td> <td>\$ 974.00</td> </tr> </tbody> </table>						CLAIMS	NUMBER FILED	NUMBER EXTRA	RATE			Total Claims	19	-20 =	0	X	\$18.00	Independent Claims	4	-3 =	1	X	\$84.00	MULTIPLE DEPENDENT CLAIMS(S) (if applicable)					\$280.00	TOTAL OF ABOVE CALCULATIONS =					\$ 974.00
CLAIMS	NUMBER FILED	NUMBER EXTRA	RATE																																
Total Claims	19	-20 =	0	X	\$18.00																														
Independent Claims	4	-3 =	1	X	\$84.00																														
MULTIPLE DEPENDENT CLAIMS(S) (if applicable)					\$280.00																														
TOTAL OF ABOVE CALCULATIONS =					\$ 974.00																														
<input checked="" type="checkbox"/> Applicant claims small entity status. See 37 CFR 1.27. The fees indicated above are reduced by 1/2. \$ 0.00																																			
SUBTOTAL = \$ 974.00																																			
Processing fee of \$130.00, for furnishing the English Translation later than <input type="checkbox"/> 20 <input type="checkbox"/> 30 months from the earliest claimed priority date (37 C.F.R. 1.492(f)). \$ 0.00																																			
TOTAL NATIONAL FEE = \$ 974.00																																			
Fee for recording the enclosed assignment (37 C.F.R. 1.21(h)). The assignment must be accompanied by an appropriate cover sheet (37 C.F.R. 3.28, 3.31). \$40.00 per property + \$ 40.00																																			
Fee for Petition to Revive Unintentionally Abandoned Application (\$1280.00 - Small Entity = \$640.00) \$ 0.00																																			
TOTAL FEES ENCLOSED = \$ 1014.00																																			
					Amount to be: refunded \$ Charged \$																														
a. <input checked="" type="checkbox"/> A check in the amount of \$1014.00 to cover the above fees is enclosed. b. <input type="checkbox"/> Please charge my Deposit Account No. 14-1140 in the amount of \$_____ to cover the above fees. A duplicate copy of this form is enclosed. c. <input checked="" type="checkbox"/> The Commissioner is hereby authorized to charge any additional fees which may be required, or credit any overpayment to Deposit Account No. 14-1140. A duplicate copy of this form is enclosed. d. <input checked="" type="checkbox"/> The entire content of the foreign application(s), referred to in this application is/are hereby incorporated by reference in this application.																																			
NOTE: Where an appropriate time limit under 37 C.F.R. 1.494 or 1.495 has not been met, a petition to revive (37 C.F.R. 1.137(a) or (b)) must be filed and granted to restore the application to pending status.																																			
SEND ALL CORRESPONDENCE TO: NIXON & VANDERHYE P.C. 1100 North Glebe Road, 8 th Floor Arlington, Virginia 22201-4714 Telephone: (703) 816-4000																																			
					SIGNATURE  Larry S. Nixon NAME																														
					25,640 February 25, 2002 REGISTRATION NUMBER Date																														

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Patent Application of

HOVELL et al

Atty. Ref.: **36-1536**

Serial No. **Unknown**

Group:

National Phase of: **PCT/GB00/03684**

International Filing Date: **25 September 2000**

Filed: **February 25, 2002**

Examiner:

For: **PACKET NETWORK INTERFACING**

* * * * *

February 25, 2002

Assistant Commissioner for Patents
Washington, DC 20231

Sir:

PRELIMINARY AMENDMENT

Prior to calculation of the filing fee and in order to place the above identified application in better condition for examination, please amend as follows:

IN THE SPECIFICATION

Page 1, after the title insert the following:

-- This application is the US national phase of international application PCT/GB00/03684 filed 25 September 2000 which designated the U.S. --.

IN THE CLAIMS

Cancel claims 19 and 20 without prejudice or disclaimer.

Please substitute the following amended claims for corresponding claims previously presented. A copy of the amended claims showing current revisions is attached.

5. (Amended) An interface as claimed in claim 1, wherein the encapsulating means comprises a plurality of different encapsulators, each encapsulator being arranged to operate in accordance with a respective encapsulation type, and the

controller is arranged to send that first type message received from the first network to the appropriate one of said plurality of different encapsulators in dependence upon the format of the destination address of that first type message.

6. (Amended) An interface as claimed in claim 3 wherein the encapsulating means comprises a plurality of different encapsulators, each encapsulator being arranged to operate in accordance with a respective encapsulation type, and the controller is arranged to send that first type message received from the first network to the appropriate one of said plurality of different encapsulators in dependence upon the format of the destination address of that first type message; and

wherein each entry of the look-up table comprises a field for containing an identifier for identifying a type of encapsulation.

7. (Amended) An interface according to claim 1, further including means for un-encapsulating an encapsulating second type message to retrieve its payload.

8. (Amended) An interface according to claim 1 wherein the protocol converter is further engaged to convert a second type message into a first type message.

12. (Amended) A method as claimed in claim 9, wherein the first predetermined format includes a first predetermined portion whose content identifies that received first type message as suitable for protocol conversion.

14. (Amended) A method as claimed in claim 9, wherein the second type address is derived directly by retrieving it from a predetermined subaddress field of the destination address.

15. (Amended) A method as claimed in claim 9, wherein the examining step comprises the substeps of retrieving the destination address from the received first type

message, and accessing a look-up table in accordance with the retrieved destination address.

18. (Amended) A method as claimed in claim 16, for use when the look-up table entries includes a second identifier field containing an identifier identifying an encapsulation type, and when there is a plurality of encapsulation types available, and including the steps of retrieving the identifier from the second identifier field of the entry having its first type address matching the destination address, and checking that the retrieved identifier is consistent with the type of encapsulation to be performed upon that received first type message.

HOVELL et al
Serial No. **Unknown**

REMARKS


Attached hereto is a marked-up version of the changes made to the claims by the current amendment. The attached page is captioned "**Version with markings to show changes made.**"

The above amendments are made to place the claims in a more traditional format.

Respectfully submitted,

NIXON & VANDERHYE P.C.

By:



Larry S. Nixon
Reg. No. 25,640

LSN:Imy

1100 North Glebe Road, 8th Floor
Arlington, VA 22201-4714
Telephone: (703) 816-4000
Facsimile: (703) 816-4100

VERSION WITH MARKINGS TO SHOW CHANGES MADE

5. (Amended) An interface as claimed in [any one of claims 1 to 4] claim 1, wherein the encapsulating means comprises a plurality of different encapsulators, each encapsulator being arranged to operate in accordance with a respective encapsulation type, and the controller is arranged to send that first type message received from the first network to the appropriate one of said plurality of different encapsulators in dependence upon the format of the destination address of that first type message.

6. (Amended) An interface as claimed in claim [5, when dependent on either claim 3 or claim 4,] 3 wherein the encapsulating means comprises a plurality of different encapsulators, each encapsulator being arranged to operate in accordance with a respective encapsulation type, and the controller is arranged to send that first type message received from the first network to the appropriate one of said plurality of different encapsulators in dependence upon the format of the destination address of that first type message; and

wherein each entry of the look-up table comprises a field for containing an identifier for identifying a type of encapsulation.

7. (Amended) An interface according to [any one of the preceding claims] claim 1, further including means for un-encapsulating an encapsulating second type message to retrieve its payload.

8. (Amended) An interface according to [any one of the preceding claims] claim 1 wherein the protocol converter is further engaged to convert a second type message into a first type message.

12. (Amended) A method as claimed in [any one of claims 9 to 10] claim 9, wherein the first predetermined format includes a first predetermined portion whose content identifies that received first type message as suitable for protocol conversion.

14. (Amended) A method as claimed in [any one of claims 9 to 13] claim 9, wherein the second type address is derived directly by retrieving it from a predetermined subaddress field of the destination address.

15. (Amended) A method as claimed in [any one of claims 9 to 14] claim 9, wherein the examining step comprises the substeps of retrieving the destination address from the received first type message, and accessing a look-up table in accordance with the retrieved destination address.

18. (Amended) A method as claimed in [either claim 16 or claim 17] claim 16, for use when the look-up table entries includes a second identifier field containing an identifier identifying an encapsulation type, and when there is a plurality of encapsulation types available, and including the steps of retrieving the identifier from the second identifier field of the entry having its first type address matching the destination address, and checking that the retrieved identifier is consistent with the type of encapsulation to be performed upon that received first type message.

PACKET NETWORK INTERFACING

This invention relates to an interface between a first network operating in accordance with a first transmission protocol and having network addresses in accordance with a first addressing convention, herein referred to as first type addresses, and a second network operating in accordance with a second transmission protocol and having network addresses in accordance with a second addressing convention, herein referred to as second type addresses; and to the tunnelling of packets across the second network from one such interface to another such interface; and relates particularly, but not exclusively, to communication between hosts in respective Internet Protocol version 6 (IPv6) domains separated by an Internet Protocol version 4 (IPv4) domain.

Herein, the terms packet and message are used interchangeably and have the same meaning, and the term Internet domain is used as a specific example of a network.

In Internet technology, it had become apparent that the initial transport protocol, IPv4, needed to be enhanced, principally to increase the available address space and add a hierarchical address structure. The result was IPv6 which has a simplified header format compared with IPv4, but which uses 128 bit addresses as compared with 32 bit addresses used in IPv4.

Readers wishing to have an overview of this general transitional area might like to access a list of Internet-Drafts, which are working documents of the Internet Engineering Task Force (IETF), at <http://www.ietf.org/1id-abstracts.txt>, and a particularly relevant document is "A Guide to the Introduction of IPv6 in the IPv4 World" <draft-ietf-ngtrans-introduction-to-ipv6-transition-01.txt>, also referred to as "Guide to IPv6 Transition".

As mentioned, the present invention relates to tunnelling. Known tunnelling techniques are of two types: configured and automatic.

A configured tunnel is created by manual configuration of a tunnelling interface between an IPv6 domain and an IPv4 domain such that all packets received from the IPv6 domain are encapsulated in IPv4 packets addressed to a specific tunnel end point, i.e. the tunnelling interface between the IPv4 domain and the remote IPv6 domain containing the destination IPv6 host.

ART 34 AMDT

5 For more detailed information on 6over4 the reader can obtain from the IETF, the document known as RFC2829, or any variant thereof, "Transmission of IPv6 over IPv4 Domains without Explicit Tunnels", by B. Carpenter and C. Jung, March 1999.

10 "Connection of IPv6 Domains via IPv4 Clouds without Explicit Tunnels", by B. Carpenter and K. Moore.

For more detailed information on Dynamic Tunnelling the reader can obtain from the IETF, the document known as draft-ietf-ngtrans-dti-00.txt.

For more detailed information on Tunnel Broker the reader can obtain from
15 the IETF, the document known as draft-ietf-ngtrans-broker-00.txt.

These known automatic tunnelling mechanisms use a variety of techniques to enable the tunnel to be automatically established:

- 6over4 Multicast
- 6to4 Special IPv6 address in which the top level aggregator (TLA) contains an identifier for the 6to4 tunnelling mechanism, and the next level aggregator (NLA) contains the IPv4 address of the tunnel end point
- Dynamic Tunnelling via DNS
- Tunnel Broker www.based.to/

European patent application EP 840 482 discloses a method of communicating
25 between an IPv4 terminal and an IPv6 terminal, in particular a method of performing
protocol conversion on IPv4 packets and IPv6 packets using converting apparatus
located between IPv4 and IPv6 networks. EP 840 482 states that there are several
known methods for communicating between IPv4 and IPv6 terminals, and discusses
the aforementioned tunnelling methods, among others. EP 840 482 asserts that there
30 are several problems with these existing methods, and accordingly provides a new
method for IPv4-IPv6 internetworking. Their method is explained by means of an
example: considering the case of an IPv4 terminal C requesting communications with
an IPv6 terminal A. The IPv4 terminal sends a name lookup request for a host A in

ART 34 AMEND

the IPv6 network. This request is received by the converting apparatus, which forwards the request to a DNS server. The DNS server returns an IPv6 network address to the converting apparatus, which assigns an IPv4 address from a pool of addresses (from a Dynamic Host Configuration Protocol (DHCP) server) to the returned IPv6 address. The converting apparatus stores the mapping between the assigned IPv4 address and the returned IPv6 address, and forwards the assigned IPv4 address to the requesting host C. In subsequent communications between hosts C and A, host C sets the destination address to be the assigned IPv4 address for outgoing packets, and sends the packets to the assigned IPv4 address of Host A; conventional IP routing ensures that the packet routes to the converting apparatus. The converting apparatus then looks up the mapping between assigned IPv4 address and IPv6 address to retrieve the IPv6 address of host A, and makes this the destination address of the packet.

EP 840 482 thus presents a new method of IPv4/IPv6 networking.

15

According to a first aspect of the present invention, there is provided, an interface for use between a first network operating in accordance with a first transmission protocol and having network addresses in accordance with a first addressing convention, herein referred to as first type addresses, and a second network operating in accordance with a second transmission protocol and having network addresses in accordance with a second addressing convention, herein referred to as second type addresses, the interface having both a first type address and a second type address and comprising:

a protocol converter arranged to convert a message having a format in accordance with the first transmission protocol, herein referred to as a first type message, into a message having a format in accordance with the second transmission protocol, herein referred to as a second type message;

means for encapsulating arranged to respond to receipt of a second type address together with a first type message by encapsulating that received first type message as the payload of a resulting encapsulating second type message, using that received second type address as the destination address of the resulting encapsulating second type message and using the second type address of the interface as the source address of the resulting encapsulating second type message;

and

an interface controller arranged to respond to receipt by the interface of a first type message from the first network by

examining the destination address of that first type message received from the first
5 network to determine whether the first type message is of a predetermined format, and if so,

sending to the protocol converter that first type message received from the first network

else, deriving, directly or indirectly, from the destination address of that first type
10 message, a second type address for use by the encapsulating means, and

sending to the encapsulating means the derived second type address together with that first type message received from the first network.

Preferably, the controller is arranged to derive the second type address directly by retrieving it from a predetermined subaddress field of the destination
15 address.

Alternatively, the controller is arranged to derive the second type address indirectly by accessing, in accordance with the destination address, a look-up table having entries in the form of a first type address associated with a second type address, and retrieving the second type address of an entry having a first type
20 address matching the destination address.

Preferably, each entry of the address conversion table comprises a field for containing an identifier for identifying whether the controller is to send that first type message received from the first network to the protocol converter or to the encapsulating means.

25 The encapsulating means may comprise a plurality of different encapsulators, each encapsulator being arranged to operate in accordance with a respective encapsulation type, and the controller being arranged to determine whether the destination address of that first type message received from the first network is of one of a respective corresponding plurality of predetermined formats, and, if so, to
30 send that first type message received from the first network to the encapsulating means corresponding to that one predetermined format.

Preferably, each entry of the address conversion table comprises a field for containing an identifier for identifying a type of encapsulation whether the controller

is to send that first type message received from the first network to the protocol converter or to the encapsulating means.

According to a second aspect of the present invention, there is provided, a method of operating an interface between a first network operating in accordance
5 with a first transmission protocol and having network addresses in accordance with a first addressing convention, herein referred to as first type addresses, and a second network operating in accordance with a second transmission protocol and having network addresses in accordance with a second addressing convention, herein referred to as second type addresses, the method comprising the steps of:

- 10 examining the destination address of a first type message received from the first network; and

if the destination address of that received first type message is of a first predetermined format, protocol converting that received first type message;

- 15 else, encapsulating that received first type message in accordance with the second transmission protocol, using, as the destination address of a resulting encapsulating second type message, a second type address derived, directly or indirectly, from the destination address of that received first type message.

- According to a third aspect of the present invention, there is provided, a method of operating an interface between a first network operating in accordance
20 with a first transmission protocol and having network addresses in accordance with a first addressing convention, herein referred to as first type addresses, and a second network operating in accordance with a second transmission protocol and having network addresses in accordance with a second addressing convention, herein referred to as second type addresses, the method comprising the steps of:

- 25 examining the destination address of a first type message received from the first network; and

if the destination address of that received first type message is of a first predetermined format, protocol converting that received first type message; and

- 30 if the destination address of that received first type message is of a second predetermined format, encapsulating that received first type message in accordance with the second transmission protocol, using, as the destination address of a resulting encapsulating second type message, a second type address derived, directly or indirectly, from the destination address of that received first type message.

ART 34 AMDT

5a

Preferably, the second predetermined address format includes an identifier identifying an encapsulation type.

According to a fourth aspect of the present invention, there is provided an interface apparatus for performing the method of the second aspect of the present invention, the interface apparatus comprising:

means for examining the destination address of a first type message received at the interface apparatus from the first network to determine whether the destination address of that received first type message is of a first predetermined format;

10 a protocol converter responsive to a determination of the examining means that the destination address of that received first type message is of the first predetermined format to convert that received first type message; and

means for encapsulating responsive to a determination of the examining means that the destination address of that received first type message is not of the first predetermined format to encapsulate that received first type message in accordance with the second transmission protocol, using, as the destination address of a resulting encapsulating second type message, a second type address derived, directly or indirectly, from the destination address of that received first type message.

20 The first predetermined format may include a first predetermined portion whose content identifies that received first type message as suitable for protocol conversion.

The first predetermined format may also include a second predetermined portion whose content constitutes the second type address used as the destination address of a resulting encapsulating second type message.

25 Preferably, the second type address is derived directly by retrieving it from a predetermined subaddress field of the destination address.

Preferably, the examining step comprises the substeps of retrieving the destination address from the received first type message, and accessing a look-up table in accordance with the retrieved destination address.

30 More preferably, in methods for use when the look-up table comprises entries in the form of a first type address associated with a second type address, retrieval of the second type address of an entry having its first type address matching the

ART 34 AMENDED

5b

destination address constitutes indirectly deriving the second type address from the destination address of that received first type message.

In methods for use when the look-up table entries include a first identifier field containing an identifier identifying whether that first type message is to be
5 protocol converted or encapsulated, there may be included the steps of retrieving the identifier from the first identifier field of the entry having its first type address matching the destination

1.0b
[1]
[2]
[3]
[4]
[5]
[6]
[7]
[8]
[9]
[10]
[11]
[12]
[13]
[14]
[15]
[16]
[17]
[18]
[19]
[20]
[21]
[22]
[23]
[24]
[25]
[26]
[27]
[28]
[29]
[30]
[31]
[32]
[33]
[34]
[35]
[36]
[37]
[38]
[39]
[40]
[41]
[42]
[43]
[44]
[45]
[46]
[47]
[48]
[49]
[50]
[51]
[52]
[53]
[54]
[55]
[56]
[57]
[58]
[59]
[60]
[61]
[62]
[63]
[64]
[65]
[66]
[67]
[68]
[69]
[70]
[71]
[72]
[73]
[74]
[75]
[76]
[77]
[78]
[79]
[80]
[81]
[82]
[83]
[84]
[85]
[86]
[87]
[88]
[89]
[90]
[91]
[92]
[93]
[94]
[95]
[96]
[97]
[98]
[99]
[100]

The first predetermined format may also include a second predetermined portion whose content constitutes the second type address used as the destination address of a resulting encapsulating second type message.

- Preferably, the second type address is derived directly by retrieving it from a
5 predetermined subaddress field of the destination address.

Preferably, the examining step comprises the substeps of retrieving the destination address from the received first type message, and accessing a look-up table in accordance with the retrieved destination address.

- More preferably, in methods for use when the look-up table comprises entries
10 in the form of a first type address associated with a second type address, retrieval of the second type address of an entry having its first type address matching the destination address constitutes indirectly deriving the second type address from the destination address of that received first type message.

- In methods for use when the look-up table entries include a first identifier
15 field containing an identifier identifying whether that first type message is to be protocol converted or encapsulated, there may be included the steps of retrieving the identifier from the first identifier field of the entry having its first type address matching the destination address, and checking that the retrieved identifier is consistent with whichever of protocol conversion or encapsulation is to be performed
20 upon that received first type message.

- In methods for use when the look-up table entries include a second identifier field containing an identifier identifying an encapsulation type, and when there is a plurality of encapsulation types available, there may be included the steps of retrieving the identifier from the second identifier field of the entry having its first
25 type address matching the destination address, and checking that the retrieved identifier is consistent with the type of encapsulation to be performed upon that received first type message.

- Protocol converters are known for enabling an IPv6 host to send messages to an IPv4 host. When a new IPv6 host is activated in an IPv6 domain, it employs the
30 technique known as Neighbourhood Discovery (ND) to find out the identity of hosts

ART 34 AMDT

6a

with whom it can communicate directly. It broadcasts an ND message containing its IPv6 network address, and each host that receives it sends a reply message containing that host's IPv6 network address. Since the domain uses an underlying transport mechanism, say Ethernet using media access control (MAC) addresses, 5 each host receiving the ND message will retrieve the new host's IPv6 network address and also the MAC address of the new host, and the new host will retrieve from each reply message the sending host's IPv6 network address and its MAC address.

10 The new host now constructs an ND table in which each entry corresponds with a neighbouring host and comprises a first part in the form of the 128 bit IPv6 address of that neighbouring host, and a second part in the form of the associated MAC address.

The interface device (containing the protocol converter) between that IPv6 domain and an adjacent IPv4 domain will also have received the ND message and 15 sent a reply message, and the new host will have made a special Default entry in the ND table having its first part formed by 128 zeros (in variants, these are all ones) and its second part formed by the MAC address of that interface device.

Thus, when that new host wants to send a message to one of the other hosts in its domain, it constructs an IPv6 message and accesses its ND table to 20 retrieve the MAC address associated with the destination address. The message is

then encapsulated within an Ethernet packet in known manner and sent via the underlying Ethernet transport mechanism to the destination host.

If, on the other hand, the host constructs an IPv6 message having its destination address in the form of an IPv4-compatible or IPv4-mapped address, i.e. a message intended for an IPv4 host in the adjacent IPv4 domain, this destination address will not be found in the ND table. In this situation, the accessing algorithm will return the MAC address of the Default entry, and the message will be sent to the protocol converter.

Protocol converters can only convert (or translate) between corresponding fields of the headers of the IPv6 and IPv4 messages. Where a field in the header of, say, an IPv6 message has no corresponding field in the header of an IPv4 message, or vice versa, the information in that field will be lost in the protocol conversion process.

As mentioned above, tunnelling techniques are known for enabling IPv6 hosts to communicate amongst themselves when they are in spaced-apart domains. In this case, the interface device contains a tunnelling mechanism instead of a protocol converter. It will be appreciated that before now, for an IPv6 host to be able to communicate both with IPv4 hosts and with remote IPv6 hosts, it was necessary for that IPv6 host and the domain interface to be dual stack, i.e. having both IPv4 and IPv6 communication capability. If an IPv6 host was not dual stack, its accessing algorithm would return only a single MAC address for the Default entry. This would be the MAC address of the input port of a protocol converter, if the network administration had decided that the IPv6 host is to be able to communicate with IPv4 hosts, or it would be the MAC address of the input port of a tunnelling mechanism, if the network administration had decided that the IPv6 host is to be able to communicate with IPv6 hosts. That Default entry MAC address could not be a common input address for both the protocol converter and the tunnelling mechanism.

Specific embodiments of the present invention will now be described with reference to the drawings in which:

Figure 1 is a schematic diagram of an IPv4 domain interfaced with two isolated IPv6 domains;

Figure 2 is a schematic diagram of a border router;

Figure 3 is a schematic diagram of an IPv6 DNS Response message;

Figure 4 is a schematic diagram of an IPv4 DNS Response message resulting from conversion of the IPv6 DNS Response message of Figure 3;

Figure 5 is a schematic diagram of an IPv6 DNS Response message resulting from conversion of the IPv4 DNS Response message of Figure 4; and

- 5 Figure 6 is a schematic diagram showing the format of special IPv6 addresses used with the 6to4 tunnelling technique.

In Figure 1, an IPv4 domain 10 separates a first IPv6 domain 12, constituting in accordance with the present invention a first network operating in accordance with a first transmission protocol and having network addresses in accordance with a first
10 addressing convention, from a second IPv6 domain 14, constituting in accordance with the present invention a third network also operating in accordance with the first transmission protocol and having network addresses in accordance with the first addressing convention. Hosts in the IPv4 domain 10 are IPv4 only, and hosts in the IPv6 domains 12 and 14 are IPv6 only.

- 15 The IPv4 domain 10 constitutes in accordance with the present invention a second network operating in accordance with a second transmission protocol and having network addresses in accordance with a second addressing convention. For simplifying the drawings, no IPv4 hosts are shown, and in each of the IPv6 domains 12 and 14 only one IPv6 host (28 and 30, respectively, as referred to later) is shown.

- 20 The first IPv6 domain 12 is connected to the IPv4 domain 10 via a border router 16A, and the second IPv6 domain 14 is connected to the IPv4 domain 10 via a border router 16B. The border router 16B is identical to the border router 16A, and each constitutes an interface.

- 25 The IPv4 domain 10 contains a complete domain name system (DNS) 20 including a plurality of DNS servers 22, of which only two DNS servers 22A and 22B are shown, and the IPv6 domains 12 and 14 contain respective DNS servers 24 and 26.

- 30 Suppose that a host 28 in the first IPv6 domain 12 wishes to send a packet to a host 30 in the second IPv6 domain 14. Thus, for this transaction, the host 28 is referred to as the source host 28 and the host 30 is referred to as the destination host 30 .

The source host 28 knows the name of the destination host 30, so it constructs in known manner an IPv6 DNS Request message (not shown) requesting

the IPv6 address of the destination host 30 . The source host 28 sends this DNS Request message as a recursive request to its local DNS server, which in this embodiment is the DNS server 24. The DNS server 24 will, in known manner, send a number of iterative DNS Request messages (not shown) to the DNS 20 until it learns about the DNS server 26. Finally, a DNS Request message (not shown) will go to the DNS server 26 requesting the IPv6 address of the destination host 30.

As the DNS request passes from the first IPv6 domain 12 through the border router 16A to the IPv4 domain 10, it is processed by a protocol converter (PC) 32A (see Figure 2) and undergoes IPv6/IPv4 translation. Correspondingly, as the DNS request passes from the IPv4 domain 10 through the border router 16B to the second IPv6 domain 14, it is processed by a protocol converter 32B and undergoes IPv4/IPv6 translation.

The protocol converters 32A and 32B conform to a specification known as Network Address Translation-Protocol Translation (NAT-PT). They translate between IPv4 and IPv6 addresses and keep state during the time of the session, and translation between IPv4 and IPv6 DNS Request messages and DNS Response messages, including translation of IP headers and DNS payload information is controlled by an application layer gateway (ALG). In the art, an alternative term for a DNS Response message is a DNS Answer message.

For more detailed information the reader can obtain from the Internet Engineering Task Force (IETF), the document known as draft-ietf-ngtrans-natpt-05.txt, or any variant thereof, "Network Address Translation - Protocol Translation (NATPT)", by G. Tsirtsis and P. Srishuresh.

The DNS server 26 responds to the DNS Request message for the IPv6 address of the destination host 30 with a IPv6 DNS Response message 34 (see Figure 3) having a conventional format of destination address field 36, source address field 38, and response address record 40 that contains the IPv6 address of the destination host 30.

This IPv6 DNS Response message 34 travels through the second IPv6 domain 14 to the border router 16B where it becomes converted to an IPv4 DNS Response message 42 (see Figure 4) comprising destination address field 44, source address field 46, response address record 48 and, in accordance with the present invention, additional records 50 and 52, and this message 42 travels through the

IPv4 domain 10 to the border router 16A where it becomes converted to an IPv6 DNS Response message 54 comprising destination address field 56, source address field 58 and response address record 60, and sent to the source host 28. The route that the DNS Response message (in its various forms, i.e. 34, 42 and 54) takes in the domains 10, 12 and 14 depends on the DNS configuration in each of the domains, but it will have to pass through the border router 16B and the border router 16A in that order.

For convenience, the terms "field" and "record" are used synonymously and interchangeably in this description, although in the art a field is generally taken to be a component part of a record.

When the IPv6 DNS Response message 34 is received by the border router 16B via its IPv6 network interface controller 62B, it is fed in parallel to the protocol converter 32B and to a controller 64B, and also to an encapsulator 86B and a 6to4 encapsulator 90B. The controller 64B is connected via a control line 65A to control inputs of the protocol converter 32B, the encapsulator 86B and the 6to4 encapsulator 90B and by placing a suitable address on the control line 65A selects the appropriate one of these devices.

The controller 64B (i) recognises that the received message is a DNS Response message and enables the protocol converter 32B, (ii) retrieves the IPv6 address from the response record 40 and writes this message to a storage location 66B formed by part of its internal memory 68B; (iii) writes the IPv4 address of the border router 16B, i.e. the IPv4 address of the tunnel terminating endpoint on the border router 16B, to a storage location 70B also formed by part of its internal memory 68B; (iv) receives from the protocol converter 32B the converted DNS Response message 42; appends the content of the storage location 70B as the first additional record 50, and the content of the storage location 68B as the second additional record 52; and (v) sends the resulting IPv4 DNS Response message 42 to an IPv4 network interface controller 72B for transmission over the IPv4 domain 10 to the border router 16A.

In a variant, the received IPv6 DNS Response message 34 is fed only to the controller 64B, which writes this message to a storage location 74B formed by part of its internal memory 68B. The controller 64B then (i) retrieves the IPv6 address from the response record 40 and writes this message to the storage location 66B; (ii)

writes the IPv4 address of the border router 16B to the storage location 70B; (iii) generates a modified IPv6 DNS Response message by retrieving the content of the storage location 66B and appending the content of the storage location 70B as the first additional record 50, and the content of the storage location 66B as the second additional record 52; and (iv) sends this modified IPv6 DNS Response message to the protocol converter 32B. The ALG of the protocol converter 32 processes only the header and address response record of the DNS Response message to produce the resulting IPv4 DNS Response message 42, i.e. it allows the additional records to remain unchanged.

- 10 It will be appreciated that the feeding of the received message directly to the protocol converter 32B, and the sending of an enabling signal on the control line 65A to the protocol converter 32B by the controller 64B is logically equivalent to the sending of the received message to the protocol converter 32B by the controller 64B in the above variant, and constitutes sending the message to the protocol converter 32B in accordance with the present invention.

When the IPv4 DNS Response message 42 is received by the border router 16A via its IPv4 network interface controller 72A, it is fed in parallel to the protocol converter 32A and to a controller 64A.

- The controller 64A (i) receives from the protocol converter 32A the output IPv6 DNS Response message comprising destination address field 56, source address field 58, response address record 60, and additional records 50, 52; and (ii) retrieves the IPv6 address from the second additional record 52, i.e. the true IPv6 address of the destination host 30, and inserts it into the response address record 60 of that output message instead of the IPv4-compatible IPv6 address for the destination host 30, which the protocol converter 32A had generated. The controller 64A then strips off the additional records 50, 52, and sends the resulting IPv6 DNS Response message 54 (see Figure 5) to an IPv6 network interface controller 62A of the border router 16A for transmission to the source host 28.

- Additionally, the controller 64A is arranged to retrieve the IPv4 address of the tunnel terminating endpoint from the first additional record 50, to create a mapping of the IPv6 address of the destination host 30 to the address of the IPv4 tunnel terminating endpoint, and to store this mapping, i.e. create an entry, in an

IPv6/tunnel endpoint table 76A formed by part of an internal memory 68A of the controller 64A and constituting a look-up table of the present invention.

In a variant, the additional records 50, 52 are stripped from the DNS Response message prior to the retrieval of their contents. In another variant, the additional records remain attached to the DNS Response message, but this is not as efficient as stripping the additional records. In this present embodiment, each entry in the IPv6/tunnel endpoint table 76A comprises a first element 78A comprising the IPv6 address of a corresponding destination host 30, a second element 80A comprising the IPv4 address of the tunnel terminating endpoint, i.e. of the border router 16B, and third and fourth elements 82A and 84A, to be described later.

Upon receipt of the resultant IPv6 DNS Response message 54, the source host 28 retrieves the IPv6 address from its address record 60 and stores it for use in sending data packets to the destination host 30.

In known manner, the source host 28 generates for each of these data packets a header including a source address field and a destination address field, and writes the retrieved IPv6 address into the destination address field.

Upon receipt of each of these data packets at the border router 16A, the controller 64A retrieves the destination address, and, in accordance with the retrieved destination address, accesses the IPv6/tunnel endpoint table 76A. If there is a match with the contents of a first element 78A of an entry, the controller 64A retrieves the corresponding IPv4 tunnel terminating endpoint from the second element 80A of that entry, and commands an encapsulator 86A to encapsulate the packet in an IPv4 packet. Thus, the encapsulator 86A appends an IPv4 header into which it has inserted its own IPv4 address into the source field, and the retrieved IPv4 tunnel terminating endpoint address into the destination field. In this embodiment the encapsulator 86A stores its own IPv4 address for this use. In variants, the controller 64A stores this IPv4 address, and passes it, together with the retrieved IPv4 tunnel terminating endpoint address, to the encapsulator 86A when commanding it to perform encapsulation.

In this embodiment, the encapsulator 86A is arranged to receive the packet directly from the IPv6 network interface controller 62A of the border router 16A, but it does not perform encapsulation until commanded by the controller 64A. In a variant, the controller 64A receives the packet directly from the IPv6 network

interface controller 62A and passes it to the encapsulator 86A if there is a match. In practice, when the border router 16A receives a packet the controller 64A writes it into a storage location of its internal memory 68A, and the controller 64A will pass to the encapsulator 86A the address of the relevant storage location, together with

5 an instruction for the encapsulator 86A to access that storage location.

Upon receipt of the encapsulating IPv4 packet at the border router 16B, an un-encapsulator 88B of the border router 16B strips off the IPv4 header and retrieves the payload of that IPv4 packet, i.e. un-encapsulates the original IPv6 packet from the source host 28, and sends that IPv6 packet to the destination host 30. The

10 controller 64B also creates a mapping (in its IPv6/tunnel endpoint table 76B) of the IPv6 address of the source host 28 and the tunnel originating endpoint, i.e. the IPv4 address of the originating border router 16A, these being respectively retrieved from the source address field of the IPv6 header and the source address field of the IPv4 header.

When the destination host 30 returns a Reply packet, a controller 64B of the border router 16B retrieves the destination address, "IPv6 host 28", from the received Reply packet, accesses its IPv6/tunnel endpoint table 76B in accordance with the retrieved destination address (i.e. seeking a matching element 78B), retrieves the corresponding IPv4 address (element 80B), and commands an

15 encapsulator 86B to encapsulate the Reply packet in an IPv4 packet addressed to an un-encapsulator 88A of border router 16A using the IPv4 tunnel originating endpoint address just retrieved from the element 80B of the IPv6/tunnel endpoint table 76B. Upon receipt of this IPv4 packet at the border router 16A, the un-encapsulator 88A performs un-encapsulation to retrieve the Reply packet, and the border router 16A

20 then sends the retrieved Reply packet to the source host 28.

The source host 28 and the destination host 30 are now in a session in which IPv6 packets pass between them via the tunnel just established between the border routers 16A and 16B.

The above described mechanism provides for an IPv6 host, which is in an isolated IPv6 domain, to communicate with another IPv6 host, which is in another isolated IPv6 domain, via an intermediate IPv4 domain, without any knowledge of where that other IPv6 host is, and without the source IPv6 host needing to do anything different from a standard communication procedure with another IPv6 host

30

within its own IPv6 domain. The DNS server local to the source IPv6 host makes a Request via the IPv4 domain to the IPv6 DNS server that is on the same network as the destination IPv6 host, and the border routers automatically set up respective mappings associating the tunnel endpoint and the IPv6 address of the IPv6 hosts
5 behind the border routers.

In alternative embodiments, some of the entries in the IPv6/tunnel endpoint table 76A can be created by the administrative personnel of the network operator. This is known as manual configuration of a tunnel, and the tunnel is permanent until changed at a later date by the administrative personnel.

10 As shown in Figure 2, the border router 16A also comprises a 6to4 tunnelling encapsulator 90A (and 6to4 tunnelling un-encapsulator 92A) and can thus interwork with a border router which is similarly enabled, although in variants these may be omitted. The special IPv6 addresses 94 (see Figure 6) used for this technique have a three-part format of which a first part 96 having thirty two bits is a prefix
15 uniquely identifying that the packet is to be tunnelled by the 6to4 tunnelling technique, a second part 98 having thirty two bits is the IPv4 address of the 6to4 tunnel endpoint, and the third part 100 having sixty four bits is known as the interface ID which is the modified MAC address of the destination host. The second part 98 constitutes a predetermined subaddress field of the destination address of the
20 present invention.

In variants having a different tunnelling encapsulator, a different respective prefix is used for the same purpose.

In some variants of these embodiments, the controller 64A is arranged to recognise the presence of this prefix within the retrieved destination address of a
25 received packet, to retrieve from its second part 98 the IPv4 address of the 6to4 tunnel endpoint and to command the 6to4 tunnelling encapsulator 90A to handle the received packet using the retrieved IPv4 address. This constitutes deriving the second type address directly. Alternatively, the 6to4 tunnelling encapsulator 90A is arranged to retrieve the special IPv6 address and to extract from its second part 98 the IPv4
30 address of the 6to4 tunnel endpoint.

Where, as mentioned above, the controller 64A is arranged to perform prefix recognition, the prefix to be recognised is stored in a storage location of its internal

memory 68A, and this storage location can be an entry or part of an entry of the IPv6/tunnel endpoint table 76A.

The un-encapsulators 88B and 92B have respective IPv4 addresses, which are used by the corresponding encapsulators 86A and 90A in generating their
5 respective encapsulated packets.

In the preferred arrangement of border router having a plurality of different encapsulators, e.g. 86, 90, the controller 64A accesses the IPv6/tunnel endpoint table 76A in accordance with a set of match criteria to cover the range of possible situations. These are

10 (a) a tunnel has already been established by the above described DNS Request technique and a IPv6 destination-specific IPv6/IPv4 entry exists in the IPv6/tunnel endpoint table 76A;

(b) a tunnel has already been established by one of the known tunnelling techniques and an IPv6/IPv4 entry exists in the IPv6/tunnel endpoint table 76A, of
15 which the first element 78A has a first part in the form of a specific prefix corresponding to that tunnelling technique;

(c) network management personnel have manually configured the border router 16 to define a tunnel to a specific IPv6 destination host using 6to4 (or 6over4) to another border router (which might be the border router 16B or a different border
20 router (not shown) associated with a further IPv6 domain (not shown)), and for this case the IPv6/tunnel endpoint table 76A has an entry whose first element 78A is the IPv4-compatible (or IPv4-mapped) address of that destination host;

(d) network management personnel have manually configured the border router 16A to define a tunnel to unspecified IPv6 destination hosts using 6to4 (or
25 6over4) to another border router (which might be the border router 16B or a different border router associated with a further IPv6 domain), and for this case the IPv6/tunnel endpoint table 76A has an entry having its first element 78A in the form of the 6to4 (or 6over4) prefix followed by the IPv4 address of that other border router followed by null characters, and in some variants the second element 80A of
30 this entry contains null characters, whilst in yet other variants the second element 80A of this entry contains the IPv4 address of that other border router; and

(e) the table contains an entry whose first element 78A is a generic IPv4-compatible or IPv4-mapped IPv6 address, i.e. its first eighty bits are all zeros and the

final thirty two (or in a variant, forty eighty) bits are null characters (zeros), the second and fourth elements 80A and 84A contain null characters, and a third element 82A contains the identifier "PC".

The controller 64A uses its IPv6/tunnel endpoint table 76A to determine the appropriate handling of a received packet in the following manner.

If the controller 64A finds an entry having a first element 78A matching the complete retrieved destination address, then the content of the second element 80A of that entry is retrieved and used as the IPv4 destination address, i.e. that of the border router 16B, the tunnel endpoint. Additionally, the content of the third element 82A of that entry is retrieved and used as a check that the retrieved IPv4 destination address and the packet received by the border router 16A are to be processed by the encapsulator 86A. The content of the third element 82A is either an identifier for an encapsulator 86A, 90A (e.g. "EN") or an identifier for the protocol converter 32A (e.g. "PC"). As a further check, the fourth element 82A of that entry contains an identifier for the encapsulation type. In other words, for an entry matching the complete retrieved destination address, as just described, the encapsulation type identifier will be "DNS" to signify that the encapsulator 86A is to be used.

If the controller 64A finds an entry whose first element 78A has the first thirty two bits matching the first thirty two bits, i.e. the special 6to4 prefix part, of the retrieved destination address, then the third and fourth elements 82A and 84A of that entry are checked ("EN" and "6to4", respectively), the IPv4 destination address is retrieved from the second element 80A of that entry and sent with the packet received by the border router 16A to be processed by the encapsulator 90A. This constitutes indirectly retrieving the second type address from a predetermined subaddress field of the destination address.

If the retrieved destination address is either IPv4-compatible or IPv4-mapped, i.e. the packet is to be protocol converted for an IPv4 destination, then its first eighty bits will be all zeros, and the following sixteen bits will be either all zeros if the address is IPv4-compatible, or all ones if the address is IPv4-mapped. Thus the controller 64A checks to see whether its IPv6/tunnel endpoint table 76A contains an entry whose first element 78A has its first eighty bits all zeros. The content of the second element 80A of such an entry will be all null characters, because no tunnel is involved, and the content of the fourth element 84A of such an entry will be all null

characters, because no encapsulation is involved. The content of the third element 82A of that entry is retrieved and used as a check that the retrieved IPv4 destination address and the packet received by the border router 16A are to be processed by the protocol converter 32A. In this case, the content of the third element 82A is an identifier for the protocol converter 32A (e.g. "PC").

In other variants, the controller 64A is arranged to access the IPv6/tunnel endpoint table 76A, as before, and only command the 6to4 tunnelling encapsulator 90A upon detecting a match with an entry, and in this case the controller 64A passes the special IPv6 address to the 6to4 tunnelling encapsulator 90A, or alternatively the controller 64A extracts the IPv4 address of the 6to4 tunnel endpoint and passes it to the 6to4 tunnelling encapsulator 82A, or yet again the controller 64A, upon detecting such a match, retrieves the element 80A of the matching entry. This element 80A contains the IPv4 address of the 6to4 tunnel endpoint, which was inserted into that element 80A by the controller (or manually) upon creation of that entry.

In the above embodiment, the local DNS server for the source host 28 is the IPv6 DNS server 24, but in alternative embodiments it can be one of the IPv4 servers 22 of the DNS 20. In such alternative embodiments, although the host 28 can send a DNS Request message for obtaining the IPv6 address of the host 30 and, by the present invention, establish a tunnel across the IPv4 domain10, the situation is not symmetrical in that the host 30 cannot act as a source and establish a corresponding tunnel across the IPv4 domain10 to the host 28. For an IPv6 host to be contactable, i.e. act as destination, by the method of the present invention, it is necessary for the DNS server local to that IPv6 host to be an IPv6 DNS server on the same IPv6 domain as that IPv6 host, because the DNS Response message has to pass through the border router adjacent to that IPv6 host in order that the tunnel can be established. In other words, the DNS Request message has to pass through the IPv4 domain and not stop at an IPv4 DNS server acting as the local DNS server for the intended destination IPv6 host.

ART 34 AMDT

18

CLAIMS

1. An interface for use between a first network operating in accordance with a first transmission protocol and having network addresses in accordance with a first addressing convention, herein referred to as first type addresses, and a
 - 5 second network operating in accordance with a second transmission protocol and having network addresses in accordance with a second addressing convention, herein referred to as second type addresses, the interface having both a first type address and a second type address and comprising:
 - a protocol converter arranged to convert a message having a format in
 - 10 accordance with the first transmission protocol, herein referred to as a first type message, into a message having a format in accordance with the second transmission protocol, herein referred to as a second type message;
 - means for encapsulating arranged to respond to receipt of a second type address together with a first type message by encapsulating that received first type
 - 15 message as the payload of a resulting encapsulating second type message, using that received second type address as the destination address of the resulting encapsulating second type message and using the second type address of the interface as the source address of the resulting encapsulating second type message; and
 - 20 an interface controller arranged to respond to receipt by the interface of a first type message from the first network by
 - examining the destination address of that first type message received from the first network to determine whether the first type message is of a predetermined format, and if so,
 - 25 sending to the protocol converter that first type message received from the first network,
 - else,
 - deriving, directly or indirectly, from the destination address of that first type message, a second type address for use by the encapsulating means, and
 - 30 sending to the encapsulating means the derived second type address together with that first type message received from the first network.

ART 34 AMDT

2. An interface as claimed in claim 1, wherein the controller is arranged to derive the second type address directly by retrieving it from a predetermined subaddress field of the destination address.
- 5 3. An interface as claimed in claim 1, wherein the controller is arranged to derive the second type address indirectly by accessing, in accordance with the destination address, a look-up table having entries in the form of a first type address associated with a second type address, and retrieving the second type address of an entry having a first type address matching the destination address.
- 10 4. An interface as claimed in claim 3, wherein each entry of the look-up table comprises a field for containing an identifier for identifying whether the controller is to send that first type message received from the first network to the protocol converter or to the encapsulating means.
- 15 5. An interface as claimed in any one of claims 1 to 4, wherein the encapsulating means comprises a plurality of different encapsulators, each encapsulator being arranged to operate in accordance with a respective encapsulation type, and the controller is arranged to send that first type message received from the first network to the appropriate one of said plurality of different encapsulators in dependence upon the format of the destination address of that first type message.
- 20 6. An interface as claimed in claim 5, when dependent on either claim 3 or claim 4, wherein each entry of the look-up table comprises a field for containing an identifier for identifying a type of encapsulation.
- 25 7. An interface according to any one of the preceding claims, further including means for un-encapsulating an encapsulating second type message to retrieve its payload.
- 30

8. An interface according to any one of the preceding claims wherein the protocol converter is further arranged to convert a second type message into a first type message.
- 5 9. A method of operating an interface between a first network and a second network, the first network having network addresses in accordance with a first addressing convention, herein referred to as first type addresses, and transmitting messages in accordance with a first transmission protocol, herein referred to as first type messages, and the second network having network
- 10 addresses in accordance with a second addressing convention, herein referred to as second type addresses, and transmitting messages in accordance with a second transmission protocol, herein referred to as second type messages, the method comprising the steps of:
- examining the destination address of a first type message received from the
- 15 first network; and
- If the destination address of that received first type message is of a first predetermined format, protocol converting that received first type message;
- else, encapsulating that received first type message in accordance with the second transmission protocol, using, as the destination address of a resulting
- 20 encapsulating second type message, a second type address derived, directly or indirectly, from the destination address of that received first type message.
10. A method of operating an interface between a first network and a second network, the first network having network addresses in accordance with a first
- 25 addressing convention, herein referred to as first type addresses, and transmitting messages in accordance with a first transmission protocol, herein referred to as first type messages, and the second network having network addresses in accordance with a second addressing convention, herein referred to as second type addresses, and transmitting messages in accordance with a
- 30 second transmission protocol, herein referred to as second type messages, the method comprising the steps of:
- examining the destination address of a first type message received from the first network; and

ART 34 AMDT

21

if the destination address of that received first type message is of a first predetermined format, protocol converting that received first type message; and

if the destination address of that received first type message is of a second predetermined format, encapsulating that received first type message in accordance
5 with the second transmission protocol, using, as the destination address of a resulting encapsulating second type message, a second type address derived, directly or indirectly, from the destination address of that received first type message.

- 10 11. A method as claimed in claim 10, wherein the second predetermined address format includes an identifier identifying an encapsulation type.
12. A method as claimed in any one of claims 9 to 10, wherein the first predetermined format includes a first predetermined portion whose content
15 identifies that received first type message as suitable for protocol conversion.
13. A method as claimed in claim 12, wherein the first predetermined format also includes a second predetermined portion whose content constitutes the second type address used as the destination address of a resulting encapsulating
20 second type message.
14. A method as claimed in any one of claims 9 to 13, wherein the second type address is derived directly by retrieving it from a predetermined subaddress field of the destination address.
- 25 15. A method as claimed in any one of claims 9 to 14, wherein the examining step comprises the substeps of retrieving the destination address from the received first type message, and accessing a look-up table in accordance with the retrieved destination address.
- 30 16. A method as claimed in claim 15, for use when the look-up table comprises entries in the form of a first type address associated with a second type address, and wherein retrieval of the second type address of an entry having

ART 34 AMDT

its first type address matching the destination address constitutes indirectly deriving the second type address from the destination address of that received first type message.

- 5 17. A method as claimed in claim 16, for use when the look-up table entries include a first identifier field containing an identifier identifying whether that first type message is to be protocol converted or encapsulated, and including the steps of retrieving the identifier from the first identifier field of the entry having its first type address matching the destination address, and checking that the retrieved identifier is consistent with whichever of protocol conversion or encapsulation is to be performed upon that received first type message.
- 10
18. A method as claimed in either claim 16 or claim 17, for use when the look-up table entries include a second identifier field containing an identifier identifying an encapsulation type, and when there is a plurality of encapsulation types available, and including the steps of retrieving the identifier from the second identifier field of the entry having its first type address matching the destination address, and checking that the retrieved identifier is consistent with the type of encapsulation to be performed upon that received first type message.
- 15
19. An interface for use between a first network operating in accordance with a first transmission protocol and having network addresses in accordance with a first addressing convention and a second network operating in accordance with a second transmission protocol and having network addresses in accordance with a second addressing convention, the interface being substantially as herein described with reference to the drawings.
- 20
20. A method of operating an interface between a first network operating in accordance with a first transmission protocol and having network addresses in accordance with a first addressing convention and a second network operating in accordance with a second transmission protocol and having network addresses in accordance with a second addressing convention, the method being substantially as herein described with reference to the drawings.
- 25
- 30

ART 24 AMDT

21. An interface for use between a first network and a second network, the first network having network addresses in accordance with a first addressing convention, herein referred to as first type addresses, and transmitting messages in accordance with a first transmission protocol, herein referred to as first type messages, and the second network having network addresses in accordance with a second addressing convention, herein referred to as second type addresses, and transmitting messages in accordance with a second transmission protocol, herein referred to as second type messages, the interface comprising:
- means for examining the destination address of a first type message received at the interface apparatus from the first network to determine whether the destination address of that received first type message is of a first predetermined format;
- a protocol converter responsive to a determination of the examining means that the destination address of that received first type message is of the first predetermined format to convert that received first type message; and
- means for encapsulating responsive to a determination of the examining means that the destination address of that received first type message is not of the first predetermined format to encapsulate that received first type message in accordance with the second transmission protocol, using, as the destination address of a resulting encapsulating second type message, a second type address derived, directly or indirectly, from the destination address of that received first type message.

ABSTRACT
PACKET NETWORK INTERFACING

A method of interfacing, and an interface for use, between, two IPv6 domains
5 separated by an IPv4 domain. The interface comprises a protocol converter, an
encapsulator/un-encapsulator and a controller. When an IPv6 source wishes to send
to a named destination in the other IPv6 domain, the source sends a normal IPv6
address request to its local DNS server, which relays it to an IPv6 name server in the
other IPv6 domain. The response message, containing the true IPv6 address of the
10 destination is received at the remote interface, which appends to the resulting
protocol converted DNS response message a first additional record containing the
true IPv6 address, and a second additional record containing the IPv4 address of that
interface. Upon receipt at the other interface, the additional records are stripped off,
their contents stored in an entry of a table, and the true IPv6 address written into the
15 address record of the resulting IPv6 DNS response message. When the interface
receives a packet from an IPv6 host, it checks whether the destination address
matches an entry of its table, and if so sends the packet to the encapsulator together
with the IPv4 address of the remote interface. The remote interface extracts the
source address and the address of the encapsulating interface and stores these in an
20 entry in its corresponding table for use in encapsulating return packets to the source.
If, however, the destination address is recognised as being of IPv4-compatible or
IPv4-mapped format, the packet is sent to a protocol converter.

25 Figure 2.




Fig.1

2/6




Fig. 2




Fig. 3




Fig. 4




Fig. 5




Fig. 6

PACKET NETWORK INTERFACING

(Attv Dkt. No.

☒ was filed as PCT International application No. PCT/GB00/03684 on 25 September 2000
and (if applicable to U.S. or PCT application) was amended on _____

Day/Month/Year Filed
24 SEPTEMBER 1999

hereby claim the benefit under 35 U.S.C. 120/365 of all prior United States and PCT international applications listed above or below and, insofar as the subject matter of each of the claims of this application is not disclosed in such prior applications in the manner provided by the first paragraph of 35 U.S.C. 112, I acknowledge the duty to disclose material information as defined in 37 C.F.R. 1.56 which occurred between the filing date of the prior applications and the national or PCT international filing date of this application:

Status: patented
pending, abandoned

PENDING

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon. And on behalf of the owner(s) hereof, I hereby appoint NIXON & VANDERHYE P.C., 1100 North Glebe Rd., Arlington, VA 22201-4714, telephone number (703) 816-4000 (to whom all communications are to be directed), and the following attorneys-in-fact, as agents for the undersigned owner(s) in connection with the prosecution of the above-captioned application before the Patent and Trademark Office connected therewith and with the resulting patent: Arthur R. Crawford, 25301, Larry S. Nixson, 25640; Robert A. Vanderhye, 27076; James T. Hosmer, 30184; Robert W. Fairis, 31352; Richard G. Besha, 32770; Mark E. Nusbaum, 32348; Michael J. Keenan, 36406; Bryan H. Davidson, 30251; Stanley C. Spooner, 27393; Leonard C. Richmond, 28009; Duane M. Byers, 33363; Jeffrey H. Nelson, 30481; John R. Lastova, 33149; H. Warren Burnam, Jr, 29266; Thomas E. Byrne, 32205; Mary J. Wilson, 32955; J. Scott Davidson, 33489; Alan M. Kagen, 36178; Charles A. Molan, 28834; B. J. Sadock, 36653; James D. Bergquist, 32775; Udepsee S. Gill, 33734; Michael J. Shea, 34725; Donald L. Jackson, 41090; Michael J. Spill, 36653; and Steven J. Schaefer, 36653. I hereby authorize Nixon & Vanderhye to delete any attorney names/nombers no longer with the firm and to act and rely solely on instructions directly communicated from the person, assignee, attorney, firm, or other organization sending instructions to Nixon & Vanderhye on behalf of the owner(s).

2. Inventor's Signature: J. King Date: 7/11/00
Inventor: JOHN (first) R KING (last) GB (citizenship)
Residence: (city) WOODBRIIDGE (state/country) GREAT BRITAIN
Post Office Address: 2 HERTFORDS PLACE, CHILLESFORD, WOODBRIDGE, SUFFOLK
(Zip Code) IP12 3SD

FOR ADDITIONAL INVENTORS, check box ☐ and attach sheet with same information and signature and date for each.

RULE 63 (37 C.F.R. 1.63)
DECLARATION AND POWER OF ATTORNEY
FOR PATENT APPLICATION

Nixon & Vanderhye P.C. (12/95)

Page 2

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

3. Inventor's Signature: 30 S. TAU Date: 1/11/2000
Inventor: JOHN PATTERSON GB
(first) (last) (citizenship)
Residence: (city) IPSWICH (state/country) GREAT BRITAIN
Post Office Address: THE ANNEXE, 5 THE MILLS, PLAYFORD ROAD, RUSHMERE ST ANDREW, IPSWICH, SUFFOLK
(Zip Code) IP4 5RL